CLAIMS

sub A1

1. In a cluster of computing nodes having shared access to one or more volumes of data storage using a parallel file system, a method for managing the data storage, comprising:

initiating a data management (DM) application in the cluster using a data management application programming interface (DMAPI) of the parallel file system;

receiving a request submitted to the parallel file system on one of the nodes to perform an operation on a file in one of the volumes of data storage;

obtaining a data management access right from the DMAPI responsive to the request; and

performing the operation on the file using the access right.

2. A method according to 1, wherein initiating the data management application comprises creating a session of the data management application on a session node selected from among the nodes in the cluster, and wherein obtaining the data management access right comprises obtaining the right at the session node.

3. A method according to claim 2, wherein initiating the data management application comprises initiating a data migration application, so as to free storage space on at least one of the volumes of data storage, and wherein receiving the request comprises generating an event responsive to the request, and wherein obtaining the right at the session node comprises associating a DM token with the right at the session node for use in invoking a DMAPI function to be applied to the file and

associating the token with the event, and wherein performing the operation comprises migrating data at a plurality of the nodes simultaneously by presenting the token in connection with the DMAPI function.

4. A method according to claim 2, wherein receiving the request comprises receiving an invocation of a file operation submitted to the parallel file system by a user application on a source node, and sending a notification of a DM event to the session node responsive to the request, and wherein obtaining the right at the session node comprises processing the event at the session node subject to the access right.

5. A method according to claim 1, wherein obtaining the data management access right comprises acquiring a data management lock on the file, so as to restrict other data management and file operations on the file while the lock is held.

6. A method according to claim 5, wherein the operation is a data management operation, and wherein acquiring the data management lock comprises holding the lock over a sequence of multiple kernel calls in the parallel file system.

7. A method according to claim 5, wherein the operation is a file operation, and wherein acquiring the data management lock comprises holding the lock for a single kernel call in the parallel file system.

8. A method according to claim 7, wherein the file operation is one of a plurality of file operations to be performed on the file, and wherein acquiring the data management lock comprises allowing the plurality of file

operations to hold respective data management locks simultaneously without mutual conflict.

9. A method according to claim 5, wherein acquiring the data management lock comprises acquiring an exclusive lock.

10. A method according to claim 5, wherein acquiring the data management lock comprises acquiring a shared lock.

11. A method according to claim 5, wherein acquiring the data management lock comprises selecting the lock from a table of locks provided for both file operations and data management operations.

12. A method according to claim 11, wherein performing the operation comprises calling a DMAPI function to perform a data management operation, and wherein acquiring the data management lock comprises acquiring, in a course of executing the DMAPI function, one of the locks provided for the file operations for the duration of the DMAPI function, so as to enable calling the DMAPI function without presenting a DM token.

13. A method according to claim 5, wherein acquiring the data management lock comprises providing the data management lock within a hierarchy of locks supported by the parallel file system.

14. Computing apparatus, comprising:

one or more volumes of data storage, arranged to store data; and

a plurality of computing nodes, linked to access the volumes of data storage using a parallel file system, and arranged so as to enable a data management (DM) application to be initiated using a data management

application programming interface (DMAPI) of the parallel file system, such that when a request submitted to the parallel file system is received on one of the nodes to perform an operation on a file in one of the volumes of data storage, a data management access right is obtained from the DMAPI responsive to the request, and the operation on the file is performed using the access right.

15. Apparatus according to 14, wherein the nodes are arranged to initiate the data management application by creating a session of the data management application on a session node selected from among the nodes in the cluster, and wherein the data management access right is obtained at the session node.

16. Apparatus according to claim 15, wherein the data management application comprises a data migration application, which frees storage space on at least one of the volumes of data storage, and wherein an event is generated responsive to the request, causing the session node to associate a DM token with the right for use in invoking a DMAPI function to be applied to the file and to associate the token with the event, and wherein data are migrated at the plurality of the nodes simultaneously by presenting the token in connection with the DMAPI function.

17. Apparatus according to claim 15, wherein the request comprises an invocation of a file operation submitted to the parallel file system by a user application on a source node, and wherein the nodes are arranged so that a notification of a DM event is sent to the session node responsive to the request, and wherein the event is

processed at the session node subject to the access right.

18. Apparatus according to claim 14, wherein the data management access right is obtained by acquiring a data management lock on the file, so as to restrict other data management and file operations on the file while the lock is held.

19. Apparatus according to claim 18, wherein the operation is a data management operation, and wherein the data management lock is held over a sequence of multiple kernel calls in the parallel file system.

20. Apparatus according to claim 18, wherein the operation is a file operation, and wherein the data management lock is held for a single kernel call in the parallel file system.

21. Apparatus according to claim 20, wherein the file operation is one of a plurality of file operations to be performed on the file, and wherein the plurality of file operations are allowed to hold respective data management locks simultaneously without mutual conflict.

22. Apparatus according to claim 18, wherein the data management lock comprises an exclusive lock.

23. Apparatus according to claim 18, wherein the data management lock comprises a shared lock.

24. Apparatus according to claim 18, wherein the data management lock is selected from a table of locks provided for both file operations and data management operations.

25. Apparatus according to claim 24, wherein the operation comprises a DMAPI function called to perform a data management operation, and wherein the data management lock comprises one of the locks provided for the file operations, which is acquired, in a course of executing the DMAPI function, for the duration of the DMAPI function, so as to enable calling the DMAPI function without presenting a DM token.

26. Apparatus according to claim 18, wherein the data management lock is provided within a hierarchy of locks supported by the parallel file system.

27. A computer software product providing a data management application programming interface (DMAPI) for use in a cluster of computing nodes having shared access to one or more volumes of data storage using a parallel file system, the product comprising a computer-readable medium in which program instructions are stored, which instructions, when read by the computing nodes, cause a data management (DM) application to be initiated using the DMAPI, such that when a request submitted to the parallel file system is received on one of the nodes to perform an operation on a file in one of the volumes of data storage, a data management access right is obtained from the DMAPI responsive to the request, and the operation on the file is performed using the access right.

28. A product according to claim 27, wherein the instructions cause the data management application to be initiated by creating a session of the data management application on a session node selected from among the

nodes in the cluster, and wherein the data management access right is obtained at the session node.

29. A product according to claim 28, wherein the data management application comprises a data migration application, which frees storage space on at least one of the volumes of data storage, and wherein the instructions cause an event to be generated responsive to the request, causing the session node to associate a DM token with the right for use in invoking a DMAPI function to be applied to the file and to associate the token with the event, and wherein data are migrated at the plurality of the nodes simultaneously by presenting the token in connection with the DMAPI function.

30. A product according to claim 28, wherein the request comprises an invocation of a file operation submitted to the parallel file system by a user application on a source node, and wherein the instructions cause a notification of a DM event to be sent to the session node responsive to the request and cause the event to be processed at the session node subject to the access right.

31. A product according to claim 27, wherein the data management access right is obtained by acquiring a data management lock on the file, so as to restrict other data management and file operations on the file while the lock is held.

32. A product according to claim 31, wherein the operation is a data management operation, and wherein the data management lock is held over a sequence of multiple kernel calls in the parallel file system.

33. A product according to claim 31, wherein the operation is a file operation, and wherein the data management lock is held for a single kernel call in the parallel file system.

34. A product according to claim 33, wherein the file operation is one of a plurality of file operations to be performed on the file, and wherein the plurality of file operations are allowed to hold respective data management locks simultaneously without mutual conflict.

35. A product according to claim 31, wherein the data management lock comprises an exclusive lock.

36. A product according to claim 31, wherein the data management lock comprises a shared lock.

37. A product according to claim 31, wherein the data management lock is selected from a table of locks provided for both file operations and data management operations.

38. A product according to claim 37, wherein the operation comprises a DMAPI function called to perform a data management operation, and wherein the data management lock comprises one of the locks provided for the file operations, which is acquired, in a course of executing the DMAPI function, for the duration of the DMAPI function, so as to enable calling the DMAPI function without presenting a DM token.

39. Apparatus according to claim 31, wherein the data management lock is provided within a hierarchy of locks supported by the parallel file system.